

IN THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently Amended) A method for incrementally scaling a file system, comprising:
adding a new file storage device to a file system having a storage space comprising at least one other file storage device having a plurality of directories and files stored thereon to form a new file system configuration; and
migrating a portion of the files from said at least one other file storage device to the new file storage device while hiding such migration from client applications that access files from the file system so as to not affect file access operations requested and performed by the client applications during the migration;
wherein the storage space of the file system is partitioned into fragments; the files in the file system are assigned to corresponding fragments; and files are assigned to corresponding fragments based on the directories the files are in.
2. (Original) The method of claim 1, wherein the portion of files that are migrated from said at least one storage device to the new storage device is selected such that the files are distributed across all of the storage devices in the file system after the migration is completed based on a relative capacity of each of the storage devices in the system.
3. (Original) The method of claim 1, wherein the file storage devices are accessed using a file system protocol, further comprising providing a storage abstraction layer between the client applications and the file system protocol, said storage abstraction layer providing an interface to the client applications that presents the file system as a virtual file system.
4. (Original) The method of claim 3, further comprising providing information corresponding to the new file system configuration to the storage abstraction layer.
5. (Original) The method of claim 3, wherein the storage abstraction layer distributes new files created by the client applications across all of the storage devices in the file system so as to load balance access operations of the files.

6. (Original) The method of claim 3, further comprising:
filtering requests made by client applications to access a file stored on the file system,
said requests referencing a virtual storage location of the file; and
remapping the file access requests that are filtered from the virtual storage location to a
physical location on a storage device on which the file is actually stored; and
accessing the file through use of the file system protocol by referencing the physical
location of the file.
7. (Original) The method of claim 1, wherein migrating the files to the new storage
device comprises:
identifying a source location corresponding to a storage device and directory in which each file is
initially stored;
identifying a destination location for each file corresponding to a directory on the new
storage device the file is to be stored in;
copying each file from the source location to the destination location;
deleting each file from its source location;
monitoring for any file access requests made by an client application while the file is
being migrated; and
aborting the migration of the file if a file access request is made during the migration of
the file.
8. (Previously Presented) The method of claim 7, further comprising putting any file
whose migration is aborted into a queue such that the migration of such file may be retried at a
future time.
9. (Original) The method of claim 7, further comprising:
providing a lock on each file during its migration; and
allowing the lock to be stolen by a client application if the client application requests
access to the file during its migration.
10. (Original) The method of claim 7, further comprising:
providing a lock token for each file opened by a client application, said token identifying
that its corresponding file is currently in use and not available to be migrated.

11. (Original) The method of claim 7, when each token is assigned an expiration time after which the token is no longer valid.
12. (Original) The method of claim 11, further comprising:
putting a file having an unexpired token into a queue such that the migration for such file may be retried at a future time; and
migrating the file after the token has expired.
13. (Canceled)
14. (Canceled)
15. (Currently Amended) The method of claim [[13]] 1, wherein the directories are assigned to corresponding fragments in a random manner.
16. (Currently Amended) The method of claim [[13]] 1, further comprising selecting (a) set(s) of fragments to be migrated when a new storage device is added to the system.
17. (Original) The method of claim 16, wherein the set(s) of fragments that are selected comprise a portion of a total number of directories on all of the storage devices in the file system such that after the set of fragments are migrated, each storage device has a proportionate amount of directories based upon its relative capacity.
18. (Original) The method of claim 1, further comprising:
providing an administrative tool that enables a user to add a new storage device to the configuration of the file system; and
automatically selecting the portion of files to be migrated to the new storage device based on the new configuration.
19. (Original) The method of claim 1, wherein the file system comprises a virtual volume corresponding to storage space provided by at least one storage device, said virtual volume including a plurality of virtual directories in which virtual files may be stored and having configuration data stored on the file system that maps virtual directories to physical directories.
20. (Original) The method of claim 19, wherein the configuration information comprises a master directory stored on a storage device, said master directory including a plurality of subdirectories, each corresponding to a respective virtual directory and having an encoded

pointer that points to a location on the file system where files corresponding to the virtual directory are physically stored.

21. (Original) The method of claim 20, wherein the configuration information further comprises a fragment map that identifies what storage device a directory and its files are stored on based upon the fragment(s) the directory is assigned to.

22. (Currently Amended) A method for load balancing file access on a network file system having a storage space provided by a plurality of network storage devices in which a plurality of files are stored, comprising:

partitioning the storage space into a plurality of fragments, each fragment being mapped to one of said plurality of network storage devices;

assigning files among said plurality of files to fragments such that each fragment, on average, comprises a substantially equal number of files;

migrating files among said plurality of files from network storage devices on which they are initially stored to other network storage devices corresponding to the fragment they are assigned to in a manner such that the migration of files are undetectable to client applications that access the network file system;

wherein each file is assigned to its corresponding fragment based upon the directory the file resides in.

23. (Previously Presented) The method of claim 22, further comprising assigning new files that are created by the client applications to fragments on a random basis.

24. (Canceled)

25. (Currently Amended) The method of claim [[24]] 22, further comprising splitting directories into a plurality of portions, wherein each directory portion of files is assigned to a respective fragment.

26. (Original) The method of claim 22, further comprising providing a storage abstraction layer that enables the client applications to access the network file system as a virtual storage space including at least one virtual volume comprising a plurality of virtual directories and file names.

27. (Original) The method of claim 26, further comprising providing the storage abstraction layer with access to a fragment map that maps each fragment to a storage device to which the fragment is hosted.
28. (Original) The method of 27, wherein each virtual directory has a corresponding physical directory on one of said plurality of network storage devices, and wherein each virtual volume includes data stored on a network storage device that links each virtual directory to its corresponding physical directory.
29. (Original) The method of claim 28, wherein the data that links the virtual and physical directories comprises a master directory that includes a plurality of subdirectories stored on a network storage device, each subdirectory being named based on a corresponding virtual directory name and including at least one file having a name comprising indicia that identifies the location of the physical directory on the network file system corresponding to the virtual directory.
30. (Original) The method of claim 29, wherein said indicia pointer comprises a first portion that identifies the fragment the files are assigned to and a second portion identifying a name of the physical directory in which the files are stored.
- 31 – 40. (Canceled)